
Kernel Mirror Prox and RKHS Gradient Flow for Mixed Functional Nash Equilibrium

Pavel Dvurechensky¹ Jia-Jie Zhu¹

Abstract

The theoretical analysis of machine learning algorithms, such as deep generative modeling, motivates multiple recent works on the Mixed Nash Equilibrium (MNE) problem. Different from MNE, this paper formulates the Mixed Functional Nash Equilibrium (MFNE), which replaces one of the measure optimization problems with optimization over a class of dual functions, e.g., the reproducing kernel Hilbert space (RKHS) in the case of Mixed Kernel Nash Equilibrium (MKNE). We show that our MFNE and MKNE framework form the backbones that govern several existing machine learning algorithms, such as implicit generative models, distributionally robust optimization (DRO), and Wasserstein barycenters. To model the infinite-dimensional continuous-limit optimization dynamics, we propose the Interacting Wasserstein-Kernel Gradient Flow, which includes the RKHS flow that is much less common than the Wasserstein gradient flow but enjoys a much simpler convexity structure. Time-discretizing this gradient flow, we propose a primal-dual kernel mirror prox algorithm, which alternates between a dual step in the RKHS, and a primal step in the space of probability measures. We then provide the first unified convergence analysis of our algorithm for this class of MKNE problems, which establishes a convergence rate of $O(1/N)$ in the deterministic case and $O(1/\sqrt{N})$ in the stochastic case. As a case study, we apply our analysis to DRO, providing the first primal-dual convergence analysis for DRO with probability-metric constraints.

¹Weierstrass Institute for Applied Analysis and Stochastics, Berlin 10117, Germany. Correspondence to: <{dvureche, zhu}@wias-berlin.de>.

1. Introduction

Training state-of-the-art large-scale machine learning models typically requires stochastic optimization with non-convex objective functions, which has achieved great empirical success. However, its reliability and computational complexity must also be theoretically analyzed and understood. For example, the potential of deep generative models is remarkable in numerous machine learning applications such as image generation (Goodfellow et al., 2014), reinforcement learning (Ho & Ermon, 2016), and molecular dynamics (Noé et al., 2019). However, training generative models effectively has been a challenging topic for machine learning research that attracts both empirical and theoretical research interests.

State-of-the-art theoretical analysis of generative models formulates the training process as finding a mixed Nash Equilibrium (MNE) in a two-player zero-sum game (Hsieh et al., 2019; Domingo-Enrich et al.; Wang & Chizat, 2022; Trillos & Trillos, 2023). The MNE problem seeks the solution, i.e., MNE of the saddle-point optimization problem

$$\inf_{\mu \in \mathcal{M}} \sup_{\nu \in \mathcal{M}} F(\mu, \nu), \quad (1)$$

where $F(\mu, \nu)$ is a bi-variate objective function(al) of the probability measures μ, ν .

In machine learning, there exists a commonly-used alternative paradigm to directly optimizing the measures, such as in the MNE (1). This paradigm exploits the duality between probability measure space \mathcal{M} and a dual functional space \mathcal{F} , i.e., instead of optimizing w.r.t. a measure, we solve the *functional optimization problem*

$$\inf_{f \in \mathcal{F}} \mathcal{E}(f). \quad (2)$$

Such methods often leverage scalable learning models such as reproducing kernel Hilbert spaces and deep neural networks for parameterizing and manipulating the dual functions, instead of directly searching in the space of probability distributions as in (1). Such dual functional approaches have witnessed success in many domains of machine learning research, e.g., generative modeling (Nowozin et al.; Arjovsky et al., 2017; Gulrajani et al., 2017; Korotin et al., 2019) com-

puting optimal transport (Genevay et al., 2016), distributionally robust optimization (Zhu et al., 2021), Wasserstein barycenters (Li et al., 2020; Tiapkin et al., 2021; Korotin et al., 2021). However, optimization over those functional spaces, e.g., RKHSs or DNNs, is inherently an infinite dimensional problem, whose convergence analysis does not directly follow from the finite-dimensional setting.

As an concrete application of our functional optimization framework, we formulate the *Mixed Functional Nash Equilibrium* (MFNE), whose inner optimization is an optimization problem over a set \mathcal{F} of functions

$$\inf_{f \in \mathcal{F}} \sup_{\mu \in \mathcal{M}} F(\mu, f), \quad (3)$$

which, like the MNE, is a special case of the infinite-dimensional pure NE problem. MFNE (3) has appeared in several cutting-edge algorithms in ML research, which we discuss in detail in Section 2.1. In such cases of saddle-point optimization problems, a unified convergence analysis for those applications to the primal-dual setting is still missing. To fill this gap, this paper provides the convergence analysis for optimization problems that move in a functional space.

Centered around the MFNE (3), this paper makes the following technical contributions:

1. We model the infinite-dimensional continuous-time optimization dynamics of the general functional optimization problem (2) as RKHS gradient flows. Despite its simple structure, the RKHS gradient flow is less explored in machine learning but enjoys simpler structures and straightforward convexity instead of (generalized) geodesic convexity needed for optimization in the Wasserstein space.
2. As a non-trivial application, we model the MFNE (3) as an Interacting Wasserstein-Kernel Gradient Flow (19)-(14), which couples the now-well-known Wasserstein gradient flow and an RKHS gradient flow. We show that the time-discretization of this interacting gradient flow results in a discrete-time primal-dual kernel mirror prox algorithm for solving MFNE (3).
3. The techniques of optimizing dual (to probability measures) functions using modern learning models, such as RKHSs and DNNs, have been applied in many recent works. However, a unified convergence analysis is still missing. This paper provides a unified convergence analysis for functional optimization problems in the context of MFNE (3), with a convergence rate $\mathcal{O}(\frac{1}{N})$ via a primal-dual kernel mirror prox algorithm. To the best of our knowledge, it is the first analysis with the kernel mirror prox steps in the dual functional space, which differs from the typical mean-field analysis of measure optimization.
4. As a case study, we apply our analysis to distributionally robust optimization (DRO) to establish a convergence rate via primal-dual stochastic kernel mirror prox in the case of maximum mean discrepancy ambiguity sets. To the best of our knowledge, this is the first primal-dual continuous optimization guarantee for DRO.
5. Last but not least, the unification perspective of a few learning tasks, provided by MFNE and MKNE, highlights our Kernel Mirror Prox as a general-purpose functional optimization algorithm for optimizing over probability measures in the dual space, with theoretical convergence guarantee. This is similar to general-purpose algorithms such as Langevin Monte-Carlo, Stein variational gradient descent.

2. Preliminaries

Notation We use $\mathcal{M} = \text{Prob}(\bar{\Omega})$ to denote the space of probability measures defined on the closure of a closed bounded convex domain $\Omega \subset \mathbb{R}^d$. We say that (3) is convex-concave if the inner problem is concave maximization and the outer problem is convex minimization. The convexity notion, if not otherwise specified, refers to the regular convexity notion (defined in (13)). For PDE gradient flows, the states are functions of both time and space, for example, $u(t, x)$. When there is no ambiguity, we write $u(t) := u(t, \cdot)$ to denote the function at evolutionary time t . If not otherwise specified, proofs are deferred to the appendix.

2.1. Duality of metrics on probability measures

One reason behind the ubiquity of the functional optimization problem (2) and MFNE (3) in machine learning problems is that common probability metrics admit a dual characterization. For example, the *optimal transport* distance (Santambrogio, 2015; Ambrosio et al., 2008), e.g., p -Wasserstein metric W_p , can be characterized in the dual space via the dual Kantorovich problem

$$\begin{aligned} \text{OT}(\mu, \nu) &= \sup_{\psi_1, \psi_2} \int \psi_1 d\mu + \int \psi_2 d\nu \\ &\text{s.t. } \psi_1(x) + \psi_2(y) \leq c(x, y), \forall x, y, \text{ a.e.} \end{aligned} \quad (4)$$

which is an infinite-dimensional optimization problem with an infinite constraint. Bounded continuous functions ψ_i are referred to as the Kantorovich potential functions. $c(x, y)$ is the transport cost function associated with the transport. In the machine learning literature, researchers have explored this dual formulation by directly parameterizing the Kantorovich potential, e.g., using RKHS functions (Genevay et al., 2016) or ICNNs (Li et al., 2020; Korotin et al., 2019).

Another commonly used metric is the integral probability

metric (IPM), which is defined via the weak norm formulation

$$\text{IPM}(\mu, \nu) = \sup_{f \in \mathcal{F}} \int f d(\mu - \nu).$$

One particular choice of the test function family is the RKHS-norm-ball $\mathcal{F} = \{f : \|f\|_{\mathcal{H}} \leq 1\}$, which yields the kernel maximum mean discrepancy (MMD) (Gretton et al., 2012). $\mathcal{F} = \{f : \text{Lip}(f) \leq 1\}$ recovers the type-1 (Kantorovich) Wasserstein metric.

Those basic duality results characterize the relationship of the underlying geometry of the learning problem and the variational problem that optimizes w.r.t. functions, such as the IPM test functions and the Kantorovich potentials in OT. This has been exploited in several fields in machine learning, which we detail below.

Implicit generative models (IGM) We consider the following IGM formulation, such as generative moment matching networks (Li et al., 2015; Dziugaite et al.) and generative adversarial networks (GAN) (Goodfellow et al., 2014),

$$\inf_{G_\theta} \mathbb{E}_Z \mathcal{D}(P, G_\theta(Z)), \quad (6)$$

where \mathcal{D} can be chosen as a discrepancy measure such as the f -divergence family, optimal transport distance, or kernel maximum mean discrepancy (MMD). In IGM, P is often taken to be the training data distribution. Recent theoretical analysis of the optimization for training (6) lifts the non-convex optimization problem to the space of probability measures, which is an instantiation of a mixed Nash equilibrium (MNE) problem (1). The problem can then be cast into this paper's general MFNE formulation (3), for example, by choosing the function \mathcal{D} as the integral probability metric (IPM) family

$$\inf_{\mu \in \mathcal{M}} \sup_{f \in \mathcal{F}} \left\{ \int f(x) dP(x) - \mathbb{E}_{\theta \sim \mu} \int f(g_\theta(z)) dQ(z) \right\}, \quad (7)$$

where $g_\theta(z)$ is the generator density. If the function class \mathcal{F} is chosen to be the class of 1-Lipschitz functions, then the formulation is the Wasserstein GAN (Arjovsky et al., 2017; Gulrajani et al., 2017). On the other hand, if \mathcal{F} is the RKHS, this is the MMD GAN (Li et al., 2017; Bińkowski et al., 2018). Note that this lifted problem is now *convex-concave* in the optimization variables.

Distributionally robust optimization (DRO) We now consider a special case: the metric-ball-constrained DRO problem (Delage & Ye, 2010)

$$\inf_{\theta} \sup_{\mu: \mathcal{D}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu[l(\theta; x)] \quad (8)$$

Below, we consider DRO with the probability metric \mathcal{D} to be the p -Wasserstein DRO ($p \geq 1$) (Gao & Kleywegt, 2016) as well as the Kernel(-MMD) DRO (Zhu et al., 2021). Different from common DRO reformulation using conic linear duality, we propose a primal-dual reformulation of DRO.

Lemma 2.1 (Primal-dual reformulation of Wasserstein and kernel DRO). *Suppose the probability metric is chosen to the MMD, then the DRO problem (8) admits the reformulation*

$$\inf_{\theta \in \mathbb{R}^d, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}} \mathbb{E}_\mu(l - f) + \frac{1}{N} \sum_{i=1}^N f(x_i) + \epsilon \|f\|_{\mathcal{H}}. \quad (9)$$

Furthermore, it is equivalent to the smoothed optimization problem

$$\inf_{\theta \in \mathbb{R}^d, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} \left(\frac{1}{N} \sum_{i=1}^N f(x_i) + \epsilon \langle h, f \rangle + \mathbb{E}_\mu(l(\theta; x) - f(x)) \right). \quad (10)$$

Suppose the probability metric is chosen to the optimal transport metric, e.g., p -Wasserstein distance. Then, the DRO problem (8) admits the reformulation

$$\inf_{\gamma > 0, \theta \in \mathbb{R}^d, f \in \Psi_c} \sup_{\mu \in \mathcal{M}} \mathbb{E}_\mu(l - \gamma \cdot f) - \frac{\gamma}{N} \sum_{i=1}^N f^c(x_i) + \gamma \cdot \epsilon. \quad (11)$$

Ψ_c is the set of c -concave (Santambrogio, 2015) functions and $f^c(y) := \inf_x c(x, y) - f(x)$ denotes the c -transform.

Lemma 2.1 shows that those primal-dual DRO formulations have the *convex-concave* structure in μ, f as in the MFKE, and are convex in the learning model θ when the loss l is. (10) is concave in the smoothing variable h . (11) is trivially convex in the dual variable γ .

Wasserstein barycenter The Wasserstein barycenter problem (Agueh & Carlier, 2011; Li et al., 2020; Tiapkin et al., 2021; Korotin et al., 2021) can be formulated as a saddle-point optimization problem.

$$\begin{aligned} & \min_{\mu \in \mathcal{M}} \sum_{i=1}^n \alpha_i [\mathcal{W}(\mu, \nu_i)] \\ & = \min_{\mu \in \mathcal{M}} \sum_{i=1}^n \alpha_i \sup_{f_i \in \Psi_c} \left\{ \int f_i^c d\mu + \int f_i d\nu_i \right\}, \quad (12) \end{aligned}$$

where $\mu_i \in \mathcal{M}$ are given probability measures, $f_i \in \Psi_c$ are the Kantorovich potentials associated with the Wasserstein distance. f_i^c again denote the c -transform.

Note that in the aforementioned settings of p -Wasserstein ($p \neq 1$) metric, one may parameterize Kantorovich potential functions using tools such as the random Fourier features (Rahimi & Recht, 2007) and input convex neural networks (Amos et al., 2017), such as done in (Genevay et al., 2016; Makkuva et al., 2020; Korotin et al., 2019). In our theoretical analysis, we focus on the RKHS functions (and hence the kernel MMD) setting in the rest of the paper due to the difficulty in theoretically characterizing the approximation error of deep models such as ICNNs.

2.2. Gradient Flow in the Wasserstein and Hilbert Spaces

Recent theoretical analysis of generative models via MNE, e.g., (Hsieh et al., 2019; Domingo-Enrich et al.; Wang & Chizat, 2022; Trillos & Trillos, 2023), adopted the mean-field limit point of view closely related to the mathematical topic on PDE gradient flow of probability measures (Ambrosio et al., 2008; Otto, a; Jordan et al., 1998; Otto, b). Notably, works such as (Hsieh et al., 2019) modeled training dynamics of GAN as sampling using Langevin SDE, which is also equivalent to solving the Fokker-Planck PDE.

Intuitively, a gradient flow describes a dynamical system that is driven towards the dissipation of certain energies. This system is called a *gradient system*. For example, the dynamical system described by an ordinary differential equation in the Euclidean space that follows the negative gradient direction, $\dot{x}(t) = -\nabla f(x(t))$, $x(t) \in \mathbb{R}^d$, is a simple gradient system. One milestone of the gradient system research is the works of Otto and colleagues in deriving the Wasserstein gradient flow (Otto, a; Jordan et al., 1998; Otto, b). Rigorous characterizations of general metric gradient systems have been carried out in PDE literature, for which we refer to (Ambrosio et al., 2008; Santambrogio, 2015) for complete treatments and (Peletier, 2014; Mielke, 2023) for a first principles' introduction, whose perspective we adopt in this paper.

Recent machine learning literature has explored the *Wasserstein gradient system*, which generates the *Wasserstein gradient flow* (WGF). It describes the evolution of probability measures in the Wasserstein metric space (\mathcal{M}, W_p) , driven by some energy functionals. In the case of 2-Wasserstein metric, the metric space (\mathcal{M}, W_2) has particularly nice properties, namely, it is a so-called geodesic metric space (Ambrosio et al., 2008). The gradient flow equation of Boltzmann entropy associated with this Wasserstein gradient system is precisely the diffusion equation, which coincide with the heat equation from the PDE perspective. From the optimization perspective, those foundational works allow us to view complex dynamical systems as optimization algorithms of probability measures, which minimizes or maximizes objective function(al)s (e.g., energy) in particular

geometries (e.g., Wasserstein space). This is vividly described by a quote from Felix Otto:

The merit of the right gradient flow formulation of a dissipative evolution equation is that it separates energetics and kinetics: The energetics endow the state space with a *functional*, the kinetics endow the state space with a (Riemannian) *geometry* via the metric tensor.

While WGF has sparked significant interest within the machine learning community recently, performing optimization with WGF requires extra care since its convergence cannot be characterized using the regular notion of convexity. Recall that a functional \mathcal{E} defined on a Hilbert space \mathcal{H} is λ -convex if $\forall s \in [0, 1], \forall u_0, u_1 \in \mathcal{H}$,

$$\mathcal{E}((1-s)u_0 + su_1) \leq (1-s)\mathcal{E}(u_0) + s\mathcal{E}(u_1) - \frac{\lambda}{2}s(1-s)\|u_0 - u_1\|_{\mathcal{H}}^2. \quad (13)$$

If $\lambda > 0$, \mathcal{E} is strongly convex. This notion of convexity does not make sense in the Wasserstein space and one must summon the generalized geodesic convexity (Ambrosio et al., 2008). In contrast, gradient flows in the Hilbert space geometry enjoy a much simpler structure and stronger characterization results. We show below that the dual functional optimization problems can be characterized by a simpler gradient flow in the reproducing kernel Hilbert space, where the regular convexity notion (13) is sufficient.

3. RKHS Gradient Flow

Our starting point is to model the infinite-dimensional continuous-time optimization dynamics of the functional optimization problem (2) using the RKHS gradient flows.

$$\partial_t f + \mathcal{E}'_f(f) = 0, \quad f(0, x) = f^0(x) \in \mathcal{H}. \quad (14)$$

In principle, other function spaces in the literature can also be considered in practice, such as the random Fourier features functions, single hidden-layer neural networks (Bach, 2017), and ICNNs. Although we focus the theoretical analysis in the rest of the paper on the RKHS setting.

Equation (14) is a gradient flow equation in a reproducing kernel Hilbert space, which is much less explored in the machine learning community than the WGF above. For example, in (Arbel et al., 2019), the squared RKHS norm was used as the driving energy functional for the Wasserstein gradient flow, rather than the dissipation geometry for the flow as in our formulation. A few other works such as (Chu et al., 2020; Chewi et al., 2020; Duncan et al., 2019; Liu & Wang, 2019; Korba et al., 2021) studied the kernelized Wasserstein gradient flow in the context of Stein geometry. Those particular cases do not exploit the simplicity

of gradient flow structure in the Hilbert space. Below, we derive standard results for RKHSGF in the context of this paper for completeness. For readers unfamiliar with PDE gradient flows in the Hilbert space, we refer to (Ambrosio et al., 2008) for complete treatment and (Santambrogio, 2017; Mielke, 2023) for accessible introductions.

We first establish the main results for RKHSGF, namely, existence, uniqueness, and a powerful result known as the *evolutionary variational inequalities* (EVI) $_{\lambda}$.

Lemma 3.1 (Characterizations of RKHS gradient flow). *Suppose the energy functional \mathcal{E} is proper, upper semicontinuous, λ -convex for some $\lambda \in \mathbb{R}$ (i.e., either convex or concave), and has compact sublevel sets. Then for any initial condition in the RKHS $f(0, x) \in \mathcal{H}$, there exists a unique solution at time t , $f(t) \in \mathcal{H}$.*

Furthermore, the gradient flow solution $f(t, x)$ satisfies (EVI) $_{\lambda}$, for $t \in [0, T]$.

$$\begin{aligned} \frac{1}{2} \|f(t) - \nu\|_{\mathcal{H}}^2 &\leq \frac{1}{2} e^{-\lambda(t-s)} \|f(s) - \nu\|_{\mathcal{H}}^2 \\ &\quad + M_{\lambda}(t-s)(\mathcal{E}(\nu) - \mathcal{E}(f(t))), \\ M_{\lambda}(\tau) &= \int_0^{\tau} e^{-\lambda(\tau-s)} ds, \quad \forall \nu \in \text{dom}(\mathcal{F}) \subset \mathcal{H}. \end{aligned}$$

Using (EVI) $_{\lambda}$, we can effortlessly extract convergence results. Suppose a minimizer of the energy exists $f^* \in \inf_{f \in \mathcal{H}} \mathcal{E}(f)$, we set $\nu = f^*$, $s = 0$ in (EVI) $_{\lambda}$

$$\begin{aligned} \|f(t) - f^*\|_{\mathcal{H}}^2 &\leq e^{-\lambda t} \|f(0) - f^*\|_{\mathcal{H}}^2 \\ &\quad + 2M_{\lambda}(t-s) \left(\inf_{f \in \mathcal{H}} \mathcal{E}(f) - \mathcal{E}(f(t)) \right) \\ &\leq e^{-\lambda t} \|f(0) - f^*\|_{\mathcal{H}}^2, \end{aligned} \quad (15)$$

yielding an *exponential convergence in time* if the energy is convex (in the usual sense) w.r.t. the f variable, i.e., $\lambda > 0$. Note that the convexity condition can be further weakened using functional inequalities such as the logarithmic Sobolev inequality.

One important distinction between the functional optimization dynamics in RKHSGF and the measure optimization in MNE optimization dynamics is that we do not need advanced structures such as *generalized geodesic convexity* from the WGF setting — linearity or regular convexity is sufficient for RKHSGF. This can be seen in the following example.

Example 3.1. [Difference in convexity] In a 2-Wasserstein space (\mathcal{M}, W_2) , linear energy functional $\mathcal{E}_0(\mu) = \int V_0 d\mu$ for $\mu \in (\mathcal{M}, W_2)$, is non-convex geodesically if and only if the function $V_0(x)$ is non-convex (Ambrosio et al., 2008).

On the other hand, in an RKHS \mathcal{H} , linear energy functional $\mathcal{E}_1(f) = \langle V_1, f \rangle_{\mathcal{H}}$ for $f \in \mathcal{H}$ is always convex, regardless of the non-convexity of V_1 .

Hence, while the MNE problem "lifts" the non-convex optimization problems to the measure spaces, the resulting objective functions as in WGF are not *geodesically convex*. Hence, the resulting WGF does not necessarily converge globally even in the linear case above. In other words, there is still no free lunch despite the lifting to the linear structure. In contrast, the RKHSGF in our formulation converges under the usual convexity in the Hilbert space. This distinction of convexity has also been exploited in the context of distributionally robust optimization with nonlinear or nonconvex (in the uncertain variable) DRO objective functions using kernel methods (Zhu et al., 2021). The same does not hold true for the Wasserstein DRO, which requires Lagrangian relaxation under nonconvex objectives (Sinha et al., 2020).

The standard proof of Lemma 3.1 is via time-discretization using the following *minimizing movement scheme* (MMS), also known as the time-incremental minimization scheme

$$\hat{f}^{k+1} \in \arg \inf_{f \in \mathcal{H}} \left\{ \mathcal{E}(f) + \frac{1}{2\tau} \|f - \hat{f}^k\|_{\mathcal{H}}^2 \right\}, \quad (16)$$

$$\hat{f}^0(x) = f(0, x) \in \mathcal{H}, \quad (17)$$

where $\tau > 0$ is the step size for time-discretization. Defining the piecewise constant function $\bar{f}_{\tau}(t, x) := \hat{f}^k(x)$ for $t \in [k\tau, (k+1)\tau]$, standard PDE proofs (see, e.g., (Ambrosio et al., 2008)) guarantee that \bar{f}_{τ} converges to the continuous-time RKHSGF solution, i.e., $\bar{f}_{\tau} \rightarrow f(t)$ as $\tau \rightarrow 0$.

Therefore, a natural bi-product of the existence results for RKHSGF is the MMS step (16). This is a fully-implicit Euler discretization of the RKHSGF, which is difficult to implement in practice. In the next section, we will use the explicit step, coupled with measure-update step, to derive the *primal-dual kernel mirror prox* algorithm, which is the discrete-time counterpart to the IWKGF (19)-(14).

Example: Interacting Wasserstein-Kernel Gradient Flow

Previously, the authors of (Domingo-Enrich et al.) proposed the Interacting Wasserstein Gradient Flow as the infinite-dimensional continuous-time optimization dynamics for solving the MNE problem (1). We now propose a new coupled gradient system which alternatives between a WGF and an *RKHS gradient flow* (RKHSGF). We term the mean-field dynamics *Interacting Wasserstein-Kernel Gradient Flow* (IWKGF). Our perspective is to choose the (dual) functional space in MFNE (3) as a reproducing kernel Hilbert space (RKHS) (Wendland, 2004; Steinwart & Christmann, 2008; Berlinet & Thomas-Agnan, 2011; Schölkopf et al., 2002). That is to say, let $\mathcal{F} = \mathcal{H}$ in MFNE (3), we solve a specific variation of MFNE, which is the *Mixed Kernel Nash Equilibrium* (MKNE)

$$\inf_{f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}} F(\mu, f). \quad (18)$$

The gradient flow equations govern the IWKGF system are

$$\partial_t \mu - \nabla \cdot (\mu \cdot F'_\mu(\mu, f)) = 0, \quad \mu(0, x) = \mu^0(x) \in \mathcal{M}, \quad (19)$$

$$\partial_t f + F'_f(\mu, f) = 0, \quad f(0, x) = f^0(x) \in \mathcal{H}. \quad (20)$$

where \mathcal{H} is an RKHS. The derivatives F'_μ, F'_f are taken in the sense of the Fréchet differential. We will calculate the concrete forms in the next sections.

When viewed standalone, The properties of the RKHS gradient flow (20) has already been discussed above. The Fokker-Planck equation (19) can be viewed either as a Wasserstein gradient flow of the energy functional $\mathcal{E}_0(\mu) := F(\mu, f)$ from the PDE perspective, or as a Langevin SDE, and is already well understood in the machine learning community. We refer to (Santambrogio, 2017; Ambrosio et al., 2008) for its properties and recent works such as (Nitanda et al., 2022; Chizat & Bach, 2018; Chizat, 2022) for recent machine learning applications.

4. A Primal-Dual Kernel Mirror Prox Algorithm

To construct our mirror prox algorithm, we consider generic variable $x \in \mathbb{R}^p$ with domain \mathcal{X} and make some mild regularity assumptions. Namely, we restrict \mathcal{M} to the set of probability measures on \mathcal{X} that admit densities w.r.t. the Lebesgue measure and have a density that is continuous and positive almost everywhere on \mathcal{X} . We also assume that there is a Hilbert space \mathcal{H} and a convex and closed set $H \in \mathcal{H}$. For the sake of generality, we consider a slight variation of MKNE (18) as the following general infinite-dimensional saddle-point problem on the spaces of measures and functions

$$\inf_{f \in H \subseteq \mathcal{H}} \sup_{\mu \in \mathcal{M}} F(f, \mu). \quad (21)$$

For shortness, we denote the set of all variables by $u = (f, \mu)$. We consider here the setting of two variables only for simplicity. An extension for a more general problem formulation covering the DRO problem (10) is easy to derive. Moreover, in the next section, we consider DRO problem (10) as a particular case study and provide technical details to check that the main assumptions of this section hold for that problem. Our first main assumption in this section is as follows.

Assumption 4.1. The functional $F(f, \mu)$ is convex in f for fixed μ and concave in μ for fixed f .

4.1. Preliminaries

To construct the mirror prox algorithm for problem (21) we need, first, to introduce proximal setup, which consists of

norms, their dual, and Bregman divergences on each space of the variables.

For the space of the variable f , we use the self-dual norm of the Hilbert space $\|\cdot\|_{\mathcal{H}}$, distance generating function $d_f(f) = \frac{1}{2}\|f\|_{\mathcal{H}}^2$, which gives Bregman divergence $B_{\mathcal{H}}(f, \check{f}) = \frac{1}{2}\|f - \check{f}\|_{\mathcal{H}}^2$. This leads to the mirror step, which is an explicit version of the MMS step (16),

$$f_+ = \text{Mirr}_\eta^{f, H}(f, \xi_f) = \arg \min_{\check{f} \in H} \{ \langle \check{f}, \eta \xi_f \rangle + \frac{1}{2} \|\check{f} - f\|_{\mathcal{H}}^2 \}. \quad (22)$$

For the space of the variable μ , we follow (Hsieh et al., 2019) and, first, introduce the Total Variation norm for the elements of \mathcal{M} $\|\mu\|_{TV} = \sup_{\|\xi\|_{L^\infty} \leq 1} \int \xi d\mu = \sup_{\|\xi\|_{L^\infty} \leq 1} \langle \xi, \mu \rangle$, where $\|\xi\|_{L^\infty}$ is the L^∞ -norm of functions. To define the mirror step, we use (negative) Shannon entropy and its Fenchel dual defined respectively as

$$\Phi(\mu) = \int d\mu \ln \frac{d\mu}{dx}, \quad \Phi^*(\xi) = \ln \int e^\xi dx \quad (23)$$

defined for ξ from the space \mathcal{F} of all bounded integrable functions on \mathcal{X} . The corresponding Bregman divergence is the relative entropy given by

$$D_\Phi(\mu, \check{\mu}) = \int d\mu \ln \frac{d\mu}{d\check{\mu}}. \quad (24)$$

This leads to the mirror step (Hsieh et al., 2019) [Theorem 1]

$$\begin{aligned} \mu_+ &= \text{Mirr}_\eta^\mu(\mu, \xi_\mu) = d\Phi^*(d\Phi(\mu) - \eta \xi_\mu) \\ &\equiv d\mu_+ = \frac{e^{-\eta \xi_\mu} d\mu}{\int e^{-\eta \xi_\mu} d\mu}. \end{aligned} \quad (25)$$

Our second main assumption is as follows.

Assumption 4.2. The functional $F(f, \mu)$ is Fréchet differentiable w.r.t. each variable and the derivatives are Lipschitz continuous in the following sense

$$\|F'_f(u) - F'_f(\tilde{u})\|_{\mathcal{H}} \leq L_{ff} \|f - \tilde{f}\|_{\mathcal{H}} + L_{f\mu} \|\mu - \tilde{\mu}\|_{TV}, \quad (26)$$

$$\|F'_\mu(u) - F'_\mu(\tilde{u})\|_{L^\infty} \leq L_{\mu f} \|f - \tilde{f}\|_{\mathcal{H}} + L_{\mu\mu} \|\mu - \tilde{\mu}\|_{TV}. \quad (27)$$

We also denote

$$L = \max_{\kappa_1, \kappa_2 \in \{f, \mu\}} \{L_{\kappa_1 \kappa_2}\}. \quad (28)$$

4.2. Kernel Mirror Prox Algorithm and Its Analysis

The updates of the ideal general infinite-dimensional mirror prox algorithm for problem (21) are given in Algorithm 1.

Algorithm 1 Ideal General Mirror-Prox

Require: Initial guess $(\tilde{f}_0, \tilde{\mu}_0)$, step-sizes $\eta_f, \eta_\mu > 0$.

- 1: **for** $k = 0, 1, \dots, N - 1$ **do**
- 2: Compute

$$\begin{aligned} f_k &= \text{Mirr}_{\eta_f}^{f, H}(F'_f(\tilde{u}_k)), \\ \mu_k &= \text{Mirr}_{\eta_\mu}^\mu(\mu_k, -F'_\mu(\tilde{u}_k)). \end{aligned}$$

- 3: Compute

$$\begin{aligned} \tilde{f}_{k+1} &= \text{Mirr}_{\eta_f}^{f, H}(F'_f(u_k)), \\ \tilde{\mu}_{k+1} &= \text{Mirr}_{\eta_\mu}^\mu(\mu_k, -F'_\mu(u_k)). \end{aligned}$$

- 4: **end for**

- 5: Compute $\bar{u}_N = \frac{1}{N} \sum_{k=0}^{N-1} u_k$.

For the analysis of the mirror prox algorithm we need the following auxiliary results. The first one is used for the mirror steps applied to the variable f .

Lemma 4.3. *Let \mathcal{H} be (possibly finite-dimensional) Hilbert space and let $H \subset \mathcal{H}$ be convex and closed. Let $\tilde{h} \in H$ and $\xi, \tilde{\xi} \in \mathcal{H}^* = \mathcal{H}$, and*

$$h = \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \xi \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_{\eta}^{h, H}(\tilde{h}, \xi), \quad (29)$$

$$\tilde{h}_+ = \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \tilde{\xi} \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_{\eta}^{h, H}(\tilde{h}, \tilde{\xi}). \quad (30)$$

Then, for any $\hat{h} \in H$

$$\begin{aligned} \langle h - \hat{h}, \eta \tilde{\xi} \rangle &\leq \frac{1}{2} \|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 \\ &\quad + \frac{\eta^2}{2} \|\tilde{\xi} - \xi\|_{\mathcal{H}}^2 - \frac{1}{2} \|h - \tilde{h}\|_{\mathcal{H}}^2. \end{aligned}$$

The second result characterizes the mirror step with respect to the measure μ .

Lemma 4.4 ((Hsieh et al., 2019) [Lemma 5]). *Let $\tilde{\mu} \in \mathcal{M}$ and $\xi, \tilde{\xi} \in \mathcal{F}$, and*

$$\mu = \text{Mirr}_{\eta}^\mu(\tilde{\mu}, \xi), \quad \tilde{\mu}_+ = \text{Mirr}_{\eta}^\mu(\tilde{\mu}, \tilde{\xi}). \quad (31)$$

Then, for any $\hat{\mu} \in \mathcal{M}$

$$\begin{aligned} \langle \mu - \hat{\mu}, \eta \tilde{\xi} \rangle &\leq D_\Phi(\hat{\mu}, \tilde{\mu}) - D_\Phi(\hat{\mu}, \tilde{\mu}_+) \\ &\quad + \frac{\eta^2}{8} \|\tilde{\xi} - \xi\|_{L^\infty}^2 - 2\|\mu - \tilde{\mu}\|_{TV}^2. \end{aligned} \quad (32)$$

The following result gives the convergence rate of Algorithm 1.

Theorem 4.5. *Let Assumptions 4.1, 4.2 hold. Let also the stepsizes in Algorithm 1 satisfy $\eta_f = \eta_\mu = \frac{1}{16L}$, where L is defined in (28). Then, for any compact set $U = U_f \times U_\mu \subseteq H \times \mathcal{M}$, the sequence $(\tilde{f}_N, \tilde{\mu}_N)$ generated by Algorithm 1 satisfies*

$$\begin{aligned} &\max_{\mu \in U_\mu} F(\tilde{f}_N, \mu) - \min_{f \in U_f} F(f, \tilde{\mu}_N) \\ &\leq \frac{8L}{N} \max_{u \in U} \left(\|f - \tilde{f}_0\|_{\mathcal{H}}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) \right). \end{aligned}$$

4.3. Analysis of Stochastic Kernel Mirror Prox

To account for potential inexactness in the first-order information, we assume that instead of exact derivatives, the algorithm uses their inexact counterparts $\tilde{F}'_f(u)$, $\tilde{F}'_\mu(u)$ that may be random and are assumed to satisfy the following assumption.

Assumption 4.6.

$$\begin{aligned} F'_f(u) &= \mathbb{E} \tilde{F}'_f(u), \quad F'_\mu(u) = \mathbb{E} \tilde{F}'_\mu(u), \quad (33) \\ \mathbb{E} \|F'_f(u) - \tilde{F}'_f(u)\|_{\mathcal{H}}^2 &\leq \sigma_f^2, \quad \mathbb{E} \|F'_\mu(u) - \tilde{F}'_\mu(u)\|_{L^\infty}^2 \leq \sigma_\mu^2. \end{aligned} \quad (34)$$

Theorem 4.7. *Let Assumptions 4.1–4.6 hold. Let also in Algorithm 1 the stochastic derivatives be used instead of the deterministic and the stepsizes satisfy $\eta_f = \eta_\mu = \frac{1}{16L}$, where L is defined in (28). Then, for any compact set $U = H \times U_\mu \subseteq H \times \mathcal{M}$, the sequence $(\tilde{f}_N, \tilde{\mu}_N)$ generated by Algorithm 1 satisfies*

$$\begin{aligned} &\mathbb{E} \left\{ \max_{\mu \in U_\mu} F(\tilde{f}_N, \mu) - \min_{f \in U_f} F(f, \tilde{\mu}_N) \right\} \\ &\leq \frac{8L}{N} \max_{u \in U} \left(\|f - \tilde{f}_0\|_{\mathcal{H}}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) \right) + \frac{3(\sigma_f^2 + \sigma_\mu^2)}{16L}. \end{aligned}$$

Let us denote $\sigma^2 = \sigma_f^2 + \sigma_\mu^2$. As we see, Theorem 4.7 guarantees the same convergence rate as in the exact case, but up to some vicinity which is governed by the level of noise. In most cases, the σ^2/L term can be made of the same order $1/N$ by using the mini-batching technique. Indeed, a mini-batch of size N allows us to reduce the variance from σ^2 to σ^2/N . Yet, we note that in this case, N iterations will require the number of samples $O(N^2)$.

An alternative would be to use the information about the diameter of the set U . Indeed, assume that

$$\max_{u \in U} \left(\|f - \tilde{f}_0\|_{\mathcal{H}}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) \right) \leq \Omega_U^2.$$

Fixing the number of steps N and choosing $\eta_f = \eta_\mu =$

$\min \left\{ \frac{1}{16L}, \frac{\Omega_U \sigma}{\sqrt{6N}} \right\}$, we obtain the following result

$$\mathbb{E} \left\{ \max_{\mu \in U_\mu} F(\bar{f}_N, \mu) - \min_{f \in U_f} F(f, \bar{\mu}_N) \right\} \leq \max \left\{ \frac{8L\Omega_U^2}{N}, \sqrt{\frac{3\sigma^2\Omega_U^2}{2N}} \right\}. \quad (35)$$

5. Case Study: Distributionally Robust Optimization

In this subsection we particularize the elements of Algorithm 1 for the specific DRO problem (8), (10). In contrast with the large number of reformulation techniques using the dual formulation, we propose the first principled *primal-dual convergence analysis* for DRO using our MFNE and MKNE framework.

Compared to problem (21) it has two additional variables $\theta \in \mathbb{R}^d$ and $h \in H \subset \mathcal{H}$. For these variables, the proximal setup is introduced in the same way as for the variable f . We choose \mathcal{H} to be a reproducing kernel Hilbert space with kernel k .

Our main assumptions for problem (10) are

1. l is convex w.r.t. θ for all x .
2. $L_0 = \sup_{x, \theta} \|\nabla_\theta l(\theta; x)\|_2 < +\infty$.
3. $\nabla_\theta l(\theta; x)$ is $L(x)$ -Lipschitz w.r.t. θ and $L_1 = \sup_\mu \mathbb{E}_{x \sim \mu} L(x)^2 < +\infty$.
4. $C = \sup_x k(x, x) < +\infty$.

We note that the convexity assumption w.r.t. θ is used to obtain global convergence guarantee w.r.t. θ . When it does not hold in practice, we can still execute our primal-dual kernel mirror prox for DRO. This is not possible with other existing Wasserstein or kernel DRO algorithms. Clearly, then the objective F is convex in (θ, f) for fixed (μ, h) and concave in (μ, h) for fixed (θ, f) . The Frechet derivatives of F with respect to the variables (θ, f, μ, h) are given by

$$F'_\theta = \mathbb{E}_{x \sim \mu} \nabla_\theta l(\theta; x) \quad (36)$$

$$F'_f = \int k(x, x') d\hat{\mu}(x') + \epsilon h(x) - \int k(x, x') d\mu(x') \quad (37)$$

$$= \mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h(\cdot) - \mathbb{E}_{x \sim \mu} k(\cdot, x) \quad (38)$$

$$-F'_\mu = f(\cdot) - l(\theta; \cdot) \quad (39)$$

$$-F'_h = -\epsilon f(\cdot). \quad (40)$$

Since the derivative w.r.t. θ and f have the form of expectation, we can use the following stochastic counterparts. We can take a sample of X_i 's from μ to construct an unbiased

stochastic derivative

$$\tilde{F}'_\theta = \frac{1}{N_\theta} \sum_{i=1}^{N_\theta} \nabla_\theta l(\theta; X_i). \quad (41)$$

Similarly, we can take a sample of X_i 's from μ and \hat{X}_i from $\hat{\mu}$ to construct an unbiased stochastic derivative

$$\tilde{F}'_f = \epsilon h(\cdot) + \frac{1}{N_f} \sum_{i=1}^{N_f} (k(\cdot, \hat{X}_i) + k(\cdot, X_i)). \quad (42)$$

We summarize the result of applying our theoretical analysis to the primal-dual DRO.

Corollary 5.1. *Assumptions 4.2, 4.6 hold for the smoothed DRO problem (10). Consequently, the results of Theorems 4.5 and 4.7 hold for the DRO problem (10).*

Therefore, we obtain $O(1/N)$ convergence rate in the deterministic case and $O(1/\sqrt{N})$ convergence rate in the stochastic case for solving DRO with kernel mirror prox.

6. Discussion

In conclusion, this paper studies the functional optimization dynamics using the continuous-time RKHS gradient flow. As a specific application, we introduce the Mixed Functional Nash Equilibrium framework that governs several learning algorithms. We model the optimization dynamics as the Interacting Wasserstein-Kernel Gradient Flow and analyze its corresponding discrete-time primal-dual kernel mirror prox algorithm. We provide the first unified convergence analysis for the MKNE problem class and the primal-dual reformulation of DRO with probability-metric constraints.

As this paper focuses on theoretical analysis, code implementation is left for future work. We also did not include the now-standard mean-field analysis using the Wasserstein gradient flow and Langevin SDE, for which we refer to recent works such as (Domingo-Enrich et al.; Chizat, 2022; Nitanda et al., 2022).

References

- Agueh, M. and Carlier, G. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2): 904–924, 2011.
- Ambrosio, L., Gigli, N., and Savare, G. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Springer Science & Business Media, 2008.
- Amos, B., Xu, L., and Kolter, J. Z. Input convex neural networks. In *International Conference on Machine Learning*, pp. 146–155. PMLR, 2017.

- Arbel, M., Korba, A., Salim, A., and Gretton, A. Maximum Mean Discrepancy Gradient Flow. *arXiv:1906.04370 [cs, stat]*, December 2019. URL <http://arxiv.org/abs/1906.04370>. arXiv: 1906.04370.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.
- Bach, F. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- Berlinet, A. and Thomas-Agnan, C. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.
- Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- Chewi, S., Le Gouic, T., Lu, C., Maunu, T., and Rigollet, P. Sgvd as a kernelized wasserstein gradient flow of the chi-squared divergence. *Advances in Neural Information Processing Systems*, 33:2098–2109, 2020.
- Chizat, L. Sparse optimization on measures with over-parameterized gradient descent. *Mathematical Programming*, 194(1-2):487–532, 2022.
- Chizat, L. and Bach, F. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.
- Chu, C., Minami, K., and Fukumizu, K. The equivalence between stein variational gradient descent and black-box variational inference. *arXiv preprint arXiv:2004.01822*, 2020.
- Dai, B., Xie, B., He, N., Liang, Y., Raj, A., Balcan, M.-F. F., and Song, L. Scalable kernel methods via doubly stochastic gradients. *Advances in neural information processing systems*, 27, 2014.
- Delage, E. and Ye, Y. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010. Publisher: INFORMS.
- Domingo-Enrich, C., Jelassi, S., Mensch, A., Rotskoff, G., and Bruna, J. A mean-field analysis of two-player zero-sum games. In *Advances in Neural Information Processing Systems*, volume 33, pp. 20215–20226. Curran Associates, Inc. URL <https://proceedings.neurips.cc/paper/2020/hash/e97c864e8ac67f7aed5ce53ec28638f5-Abstract.html>.
- Duncan, A., Nüsken, N., and Szpruch, L. On the geometry of stein variational gradient descent. *arXiv preprint arXiv:1912.00894*, 2019.
- Dziugaite, G. K., Roy, D. M., and Ghahramani, Z. Training generative neural networks via maximum mean discrepancy optimization.
- Gao, R. and Kleywegt, A. J. Distributionally Robust Stochastic Optimization with Wasserstein Distance. *arXiv:1604.02199 [math]*, July 2016. URL <http://arxiv.org/abs/1604.02199>. arXiv: 1604.02199.
- Genevay, A., Cuturi, M., Peyré, G., and Bach, F. Stochastic optimization for large-scale optimal transport. *Advances in neural information processing systems*, 29, 2016.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative Adversarial Networks. *arXiv:1406.2661 [cs, stat]*, June 2014. URL <http://arxiv.org/abs/1406.2661>. arXiv: 1406.2661.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. Publisher: JMLR. org.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved Training of Wasserstein GANs. *arXiv:1704.00028 [cs, stat]*, December 2017. URL <http://arxiv.org/abs/1704.00028>. arXiv: 1704.00028.
- Ho, J. and Ermon, S. Generative Adversarial Imitation Learning. *arXiv:1606.03476 [cs]*, June 2016. URL <http://arxiv.org/abs/1606.03476>. arXiv: 1606.03476.
- Hsieh, Y.-P., Liu, C., and Cevher, V. Finding mixed Nash equilibria of generative adversarial networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2810–2819. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/hsieh19b.html>.
- Jordan, R., Kinderlehrer, D., and Otto, F. The variational formulation of the Fokker–Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998. Publisher: SIAM.
- Korba, A., Aubin-Frankowski, P.-C., Majewski, S., and Ablin, P. Kernel Stein Discrepancy Descent. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 5719–5730. PMLR, July

2021. URL <https://proceedings.mlr.press/v139/korba21a.html>. ISSN: 2640-3498.
- Korotin, A., Egiazarian, V., Asadulaev, A., Safin, A., and Burnaev, E. Wasserstein-2 generative networks. *arXiv preprint arXiv:1909.13082*, 2019.
- Korotin, A., Li, L., Solomon, J., and Burnaev, E. Continuous Wasserstein-2 Barycenter Estimation without Minimax Optimization. *arXiv:2102.01752 [cs, stat]*, February 2021. URL <http://arxiv.org/abs/2102.01752>. arXiv: 2102.01752.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. Mmd gan: Towards deeper understanding of moment matching network. *Advances in neural information processing systems*, 30, 2017.
- Li, L., Genevay, A., Yurochkin, M., and Solomon, J. Continuous Regularized Wasserstein Barycenters. *arXiv:2008.12534 [cs, stat]*, October 2020. URL <http://arxiv.org/abs/2008.12534>. arXiv: 2008.12534.
- Li, Y., Swersky, K., and Zemel, R. Generative moment matching networks. In *International conference on machine learning*, pp. 1718–1727. PMLR, 2015.
- Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose bayesian inference algorithm. *Advances in neural information processing systems*, 29, 2016.
- Liu, Q. and Wang, D. Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. *arXiv:1608.04471 [cs, stat]*, September 2019. URL <http://arxiv.org/abs/1608.04471>. arXiv: 1608.04471.
- Makkuva, A., Taghvaei, A., Oh, S., and Lee, J. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pp. 6672–6681. PMLR, 2020.
- Mielke, A. An introduction to the analysis of gradients systems, June 2023. URL <http://arxiv.org/abs/2306.05026>. arXiv:2306.05026 [math-ph].
- Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust Stochastic Approximation Approach to Stochastic Programming. *SIAM Journal on Optimization*, 19(4):1574–1609, January 2009. ISSN 1052-6234, 1095-7189. doi: 10.1137/070704277. URL <http://epubs.siam.org/doi/10.1137/070704277>.
- Nitanda, A., Wu, D., and Suzuki, T. Convex Analysis of the Mean Field Langevin Dynamics. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pp. 9741–9757. PMLR, May 2022.
- Noé, F., Olsson, S., Köhler, J., and Wu, H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, September 2019. doi: 10.1126/science.aaw1147.
- Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc. URL <https://proceedings.neurips.cc/paper/2016/hash/cedebb6e872f539bef8c3f919874e9d7-Abstract.html>.
- Otto, F. *Double degenerate diffusion equations as steepest descent*. Citeseer, a.
- Otto, F. The geometry of dissipative evolution equations: the porous medium equation. b. Publisher: Taylor & Francis.
- Peletier, M. A. Variational Modelling: Energies, gradient flows, and large deviations. *arXiv:1402.1990 [math-ph]*, February 2014. URL <http://arxiv.org/abs/1402.1990>. arXiv: 1402.1990.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.
- Santambrogio, F. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55(58-63):94, 2015. Publisher: Springer.
- Santambrogio, F. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154, April 2017. ISSN 1664-3607, 1664-3615. doi: 10.1007/s13373-017-0101-1. URL <http://link.springer.com/10.1007/s13373-017-0101-1>.
- Schölkopf, B., Smola, A. J., Bach, F., et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Sinha, A., Namkoong, H., Volpi, R., and Duchi, J. Certifying Some Distributional Robustness with Principled Adversarial Training. *arXiv:1710.10571 [cs, stat]*, May 2020. URL <http://arxiv.org/abs/1710.10571>. arXiv: 1710.10571.
- Steinwart, I. and Christmann, A. *Support vector machines*. Springer Science & Business Media, 2008.
- Tiapkin, D., Gasnikov, A., and Dvurechensky, P. Stochastic Saddle-Point Optimization for Wasserstein Barycenters. *arXiv:2006.06763 [cs, math, stat]*, December 2021. URL <http://arxiv.org/abs/2006.06763>. arXiv: 2006.06763.

Trillos, C. G. and Trillos, N. G. On adversarial robustness and the use of wasserstein ascent-descent dynamics to enforce it. *arXiv preprint arXiv:2301.03662*, 2023.

Wang, G. and Chizat, L. An exponentially converging particle method for the mixed nash equilibrium of continuous games. *arXiv preprint arXiv:2211.01280*, 2022.

Wendland, H. *Scattered data approximation*, volume 17. Cambridge university press, 2004.

Zhu, J.-J., Jitkrittum, W., Diehl, M., and Schölkopf, B. Kernel Distributionally Robust Optimization: Generalized Duality Theorem and Stochastic Approximation. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, pp. 280–288. PMLR, March 2021. URL <https://proceedings.mlr.press/v130/zhu21a.html>. ISSN: 2640-3498.

A. Further Technical Background

A.1. Proof of Lemma 2.1

We now prove Lemma 2.1, of which a more detailed version is stated below.

Lemma A.1 (Primal-dual reformulation of Wasserstein and kernel DRO). *Suppose the probability metric is chosen to the MMD, then the DRO problem (8) admits the following equivalent primal-dual reformulations*

$$\inf_{\theta \in \mathbb{R}^d, f \in \mathcal{H}, \tau > 0} \sup_{\mu \in \mathcal{M}} \mathbb{E}_\mu(l(\theta; x) - f(x)) + \frac{1}{N} \sum_{i=1}^N f(x_i) + \frac{\tau}{2} \|f\|_{\mathcal{H}}^2 + \frac{\epsilon^2}{2\tau}, \quad (43)$$

$$\inf_{\theta \in \mathbb{R}^d, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}} \mathbb{E}_\mu(l(\theta; x) - f(x)) + \frac{1}{N} \sum_{i=1}^N f(x_i) + \epsilon \|f\|_{\mathcal{H}}. \quad (44)$$

Furthermore, it is equivalent to the smoothed optimization problem

$$\inf_{\theta \in \mathbb{R}^d, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} \left\{ \frac{1}{N} \sum_{i=1}^N f(x_i) + \epsilon \langle h, f \rangle + \mathbb{E}_\mu(l(\theta; x) - f(x)) \right\}. \quad (45)$$

Suppose the probability metric is chosen to the optimal transport metric, e.g., p -Wasserstein distance. Then, the DRO problem (8) admits the following equivalent reformulations

$$\inf_{\gamma > 0, \theta \in \mathbb{R}^d, f \in \Psi_{\gamma \cdot c}} \sup_{\mu \in \mathcal{M}} \mathbb{E}_\mu(l(\theta; x) - f(x)) - \frac{1}{N} \sum_{i=1}^N f^{\gamma \cdot c}(x_i) + \gamma \cdot \epsilon, \quad (46)$$

$$\inf_{\gamma > 0, \theta \in \mathbb{R}^d, f \in \Psi_c} \sup_{\mu \in \mathcal{M}} \mathbb{E}_\mu(l(\theta; x) - \gamma \cdot f(x)) - \frac{1}{N} \sum_{i=1}^N \gamma \cdot f^c(x_i) + \gamma \cdot \epsilon. \quad (47)$$

Ψ_c denotes the set of c -concave (Santambrogio, 2015) functions and $f^c(y) := \inf_x c(x, y) - f(x)$ denotes the c -transform.

Proof. For this proof, it suffices to consider the case where θ is fixed, since only the inner maximization is reformulated. We first prove the result for the MMD setting.

MMD setting. We consider the kernel mean embedding map as a linear constraint

$$\int \phi(x) d\mu = h, \quad (48)$$

where h is a function in \mathcal{H} . By straightforward Lagrange duality and associating the linear constraint with the multiplier in the dual space, $f \in \mathcal{H}$,

$$\inf_{\tau > 0, f \in \mathcal{H}} \sup_{\mu} \left\{ \langle l, \mu \rangle - \frac{1}{2\tau} \|h - \hat{h}\|_{\mathcal{H}}^2 + \frac{\epsilon^2}{2\tau} - \langle f, h \rangle_{\mathcal{H}} + \int \langle f, \phi(x) \rangle_{\mathcal{H}} d\mu \right\} \quad (49)$$

where \hat{h} is the kernel mean embedding of the empirical measure $\hat{P}_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$. Carrying out the quadratic optimization problem in closed form w.r.t. h and rearranging the terms, we obtain the result in (43). An optimal choice of the dual variable τ yields the equivalent reformulation (44). Smoothing via the definition of the dual norm in the Hilbert spaces, we obtain (45).

Wasserstein setting. The reformulation is a direct consequence of combining the dual Kantorovich representation of OT (4) and Lagrange duality. \square

A.2. Lemma (3.1)

Since an RKHS is a Hilbert space, Lemma (3.1) is simply the $(\text{EVI})_\lambda$ in a Hilbert space, whose proof is standard (Ambrosio et al., 2008; Santambrogio, 2015; Mielke, 2023).

A.3. Technial details on geodesic convexity in the Wasserstein space

Recent machine learning literature has explored *Wasserstein gradient flow* (WGF). The 2-Wasserstein space (\mathcal{M}, W_2) is geodesically convex, as defined below.

In a metric space $(\mathcal{M}, \mathcal{D})$ a curve $\gamma : [0, 1] \rightarrow \mathcal{M}$ is a (constant speed) geodesic if

$$\forall r, s \in [0, 1] : \quad \mathcal{D}(\gamma(r), \gamma(s)) = |s - r| \mathcal{D}(\gamma(0), \gamma(1)).$$

We refer to that as the geodesic γ connects the points $\gamma(0)$ and $\gamma(1)$ and write $\text{Geod}(\gamma(0), \gamma(1))$ for the set of all such geodesics.

Definition A.2 (Geodesic metric spaces). The metric space $(\mathcal{M}, \mathcal{D})$ is a geodesic space, if for all $u_0, u_1 \in \mathcal{M}$ there exists a geodesic connecting u_0 and u_1 .

Definition A.3 (Geodesic convexity). A functional $\mathcal{F} : \mathcal{M} \rightarrow \mathbb{R}_\infty$ is geodesically λ -convex if $\forall u_0, u_1 \in \mathcal{M} \exists \gamma \in \text{Geod}(u_0, u_1)$,

$$\mathcal{E}(\gamma(s)) \leq (1 - s)\mathcal{E}(\gamma(0)) + s\mathcal{E}(\gamma(1)) - \frac{\lambda}{2}s(1 - s)\mathcal{D}(\gamma(0), \gamma(1))^2, \quad \forall s \in [0, 1].$$

As we have seen in the main text, e.g., discussions around Example 3.1, this complication of convexity structure makes optimization in the Wasserstein more difficult than general Hilbert spaces, which motivates our approach to work in the RKHS.

A.4. Practical Consideration about Infinite-dimensional Mirror Steps

Algorithm 1 requires two infinite-dimensional mirror steps in the functional variables μ, f . The f -update, optimization w.r.t. RKHS functions, is standard in the machine learning literature, e.g., (Dai et al., 2014; Genevay et al., 2016; Tiapkin et al., 2021; Zhu et al., 2021). We now discuss the μ -update. In addition to existing works using Bregman-mirror steps, such as (Hsieh et al., 2019), which employ Langevin Monte-Carlo, we show a fully variational approach using Wasserstein gradient flows. While the PDE and SDE can describe the same drift-diffusion process, the resulting optimization algorithms are different. Notably, it is possible to perform deterministic optimization steps, purely relying on the Wasserstein geometry and bypassing the discretization of Langevin SDE. The goal of the mirror step (25) is to solve the following optimization problem

$$\mu^{k+1} \in \arg \inf_{\mu \in \mathcal{M}} \int F'_\mu(\mu^k, f^k) d\mu + \frac{1}{2\tau} \mathcal{D}(\mu, \mu^k). \quad (50)$$

The optimization objective is the energy of the Fokker-Planck equation, of which we can take the time-discretization in the Wasserstein space. This results in the standard *JKO steps* for Wasserstein gradient flow (Otto, a; Jordan et al., 1998; Otto, b). Let $\xi^0 = \mu^k$, for the sub-step count $l = 0, \dots, T$, step size $s > 0$, we solve the variational problem

$$\xi^{l+1} \in \arg \inf_{\xi \in \mathcal{M}} \int F'_\mu(\mu^k, f^k) d\xi + \frac{1}{2\tau} \mathcal{D}(\xi, \mu^k) + \frac{1}{2s} W_2^2(\xi, \xi^l). \quad (51)$$

Each step can be realized, e.g., via Stein variational gradient descent (Liu & Wang, 2016).

B. Theoretical Analysis of Primal-Dual Kernel Mirror Descent in Section 4

For the sake of generality, in particular, to cover the DRO problem (10), we consider a more general saddle-point problem. Theorems 4.5 and 4.7 are obtained as corollaries of the results obtained in this section. To make this section self-contained and for the reader's convenience we repeat some definitions and results given in the main text.

We consider generic variable $x \in \mathbb{R}^p$ with domain \mathcal{X} . We denote by \mathcal{M} the set of all probability measures on \mathcal{X} that admit densities w.r.t. the Lebesgue measure and the density is continuous and positive almost everywhere on \mathcal{X} . We also assume that there are two Hilbert spaces $\mathcal{H}_f, \mathcal{H}_h$ and convex set $\Theta \subseteq \mathbb{R}^d$ and convex compact $H \in \mathcal{H}_h$. We consider the following general infinite-dimensional saddle-point problem

$$\inf_{\theta \in \Theta \subseteq \mathbb{R}^d, f(x) \in \mathcal{H}_f} \sup_{\mu \in \mathcal{M}, h(x) \in H \subseteq \mathcal{H}_h} F(\theta, f, \mu, h). \quad (52)$$

For shortness, we denote the set of all variables by $u = (\theta, f, \mu, h)$.

Our first main assumption is as follows.

Assumption B.1. The functional $F(\theta, f, \mu, h)$ is convex in (θ, f) for fixed (μ, h) and concave in (μ, h) for fixed (θ, f) .

B.1. Preliminaries

To construct the mirror prox algorithm for problem (52) we need to first introduce proximal setup, which consists of norms, their dual, and Bregman divergences on each space of the variables.

For the space of the variable θ , we introduce the standard proximal setup with the self-dual Euclidean norm $\|\cdot\|_2$, distance-generating function $d_\theta(\theta) = \frac{1}{2}\|\theta\|_2^2$, which gives Bregman divergence $B_\theta(\theta, \check{\theta}) = \frac{1}{2}\|\theta - \check{\theta}\|_2^2$. This leads to the mirror step defined as

$$\theta_+ = \text{Mirr}_\eta^{\theta, \Theta}(\theta, \xi_\theta) = \arg \min_{\check{\theta} \in \Theta} \{ \langle \check{\theta}, \eta \xi_\theta \rangle + \frac{1}{2} \|\check{\theta} - \theta\|_2^2 \}. \quad (53)$$

We note that our choice of the Euclidean proximal setup is made for simplicity and that other standard proximal setups are possible (Nemirovski et al., 2009).

For the space of the variable f , we use the self-dual norm of the Hilbert space $\|\cdot\|_{\mathcal{H}_f}$, distance generating function $d_f(f) = \frac{1}{2}\|f\|_{\mathcal{H}_f}^2$, which gives Bregman divergence $B_{\mathcal{H}_f}(f, \check{f}) = \frac{1}{2}\|f - \check{f}\|_{\mathcal{H}_f}^2$. This leads to the mirror step

$$f_+ = \text{Mirr}_\eta^{f, \mathcal{H}_f}(f, \xi_f) = \arg \min_{\check{f}} \{ \langle \check{f}, \eta \xi_f \rangle + \frac{1}{2} \|\check{f} - f\|_{\mathcal{H}_f}^2 \} = f - \eta \xi_f. \quad (54)$$

For the space of the variable μ , we follow (Hsieh et al., 2019) and, first, introduce the Total Variation norm for the elements of \mathcal{M}

$$\|\mu\|_{TV} = \sup_{\|\xi\|_{L^\infty} \leq 1} \int \xi d\mu = \sup_{\|\xi\|_{L^\infty} \leq 1} \langle \xi, \mu \rangle,$$

where $\|\xi\|_{L^\infty}$ is the L^∞ -norm of functions. To define the mirror step, we use (negative) Shannon entropy

$$\Phi(\mu) = \int d\mu \ln \frac{d\mu}{dx} \quad (55)$$

and its Fenchel dual

$$\Phi^*(\xi) = \ln \int e^{\xi} dx \quad (56)$$

defined for ξ from the space \mathcal{F} of all bounded integrable functions on \mathcal{X} . The corresponding Bregman divergence is the relative entropy given by

$$D_\Phi(\mu, \check{\mu}) = \int d\mu \ln \frac{d\mu}{d\check{\mu}}. \quad (57)$$

This leads to the mirror step (Hsieh et al., 2019) [Theorem 1]

$$\mu_+ = \text{Mirr}_\eta^\mu(\mu, \xi_\mu) = d\Phi^*(d\Phi(\mu) - \eta \xi_\mu) \equiv d\mu_+ = \frac{e^{-\eta \xi_\mu} d\mu}{\int e^{-\eta \xi_\mu} d\mu}. \quad (58)$$

Finally, for the space of the variable h , we do the same as for the variable f . Namely, we use the distance generating function $d_h(h) = \frac{1}{2}\|h\|_{\mathcal{H}_h}^2$, which gives Bregman divergence $B_{\mathcal{H}_h}(h, \check{h}) = \frac{1}{2}\|h - \check{h}\|_{\mathcal{H}_h}^2$. This leads to the mirror step

$$h_+ = \text{Mirr}_\eta^{h, H}(h, \xi_h) = \arg \min_{\check{h} \in H} \{ \langle \check{h}, \eta \xi_h \rangle + \frac{1}{2} \|\check{h} - h\|_{\mathcal{H}_h}^2 \}. \quad (59)$$

Our second main assumption is as follows

Assumption B.2. The functional $F(\theta, f, \mu, h)$ is Fréchet differentiable w.r.t. each variable and the derivatives are Lipschitz continuous in the following sense

$$\|F'_\theta(u) - F'_\theta(\tilde{u})\|_2 \leq L_{\theta\theta}\|\theta - \tilde{\theta}\|_2 + L_{\theta f}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{\theta\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{\theta h}\|h - \tilde{h}\|_{\mathcal{H}_h}, \quad (60)$$

$$\|F'_f(u) - F'_f(\tilde{u})\|_{\mathcal{H}_f} \leq L_{f\theta}\|\theta - \tilde{\theta}\|_2 + L_{ff}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{f\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{fh}\|h - \tilde{h}\|_{\mathcal{H}_h}, \quad (61)$$

$$\|F'_\mu(u) - F'_\mu(\tilde{u})\|_{L^\infty} \leq L_{\mu\theta}\|\theta - \tilde{\theta}\|_2 + L_{\mu f}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{\mu\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{\mu h}\|h - \tilde{h}\|_{\mathcal{H}_h}, \quad (62)$$

$$\|F'_h(u) - F'_h(\tilde{u})\|_{\mathcal{H}_h} \leq L_{h\theta}\|\theta - \tilde{\theta}\|_2 + L_{hf}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{h\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{hh}\|h - \tilde{h}\|_{\mathcal{H}_h}. \quad (63)$$

We also denote

$$L = \max_{\kappa_1, \kappa_2 \in \{\theta, f, \mu, h\}} \{L_{\kappa_1 \kappa_2}\}. \quad (64)$$

B.2. Mirror Prox Algorithm and Its Analysis

The updates of the general ideal infinite-dimensional Mirror Prox algorithm for problem (52) are given in Algorithm 2.

Algorithm 2 Ideal General Mirror-Prox

Require: Initial guess $(\tilde{\theta}_0, \tilde{f}_0, \tilde{\mu}_0, \tilde{h}_0)$, step-sizes $\eta_\theta, \eta_f, \eta_\mu, \eta_h > 0$.

- 1: **for** $k = 0, 1, \dots, N - 1$ **do**
- 2: Compute

$$\begin{aligned} \theta_k &= \text{Mirr}_{\eta_\theta}^{\theta, \Theta}(\tilde{\theta}_k, F'_\theta(\tilde{u}_k)), & f_k &= \text{Mirr}_{\eta_f}^{f, \mathcal{H}_f}(F'_f(\tilde{u}_k)), \\ \mu_k &= \text{Mirr}_{\eta_\mu}^\mu(\mu_k, -F'_\mu(\tilde{u}_k)), & h_k &= \text{Mirr}_{\eta_h}^{h, H}(h_k, -F'_h(\tilde{u}_k)). \end{aligned}$$

- 3: Compute

$$\begin{aligned} \tilde{\theta}_{k+1} &= \text{Mirr}_{\eta_\theta}^{\theta, \Theta}(\tilde{\theta}_k, F'_\theta(u_k)), & \tilde{f}_{k+1} &= \text{Mirr}_{\eta_f}^{f, \mathcal{H}_f}(F'_f(u_k)), \\ \tilde{\mu}_{k+1} &= \text{Mirr}_{\eta_\mu}^\mu(\mu_k, -F'_\mu(u_k)), & \tilde{h}_{k+1} &= \text{Mirr}_{\eta_h}^{h, H}(h_k, -F'_h(u_k)). \end{aligned}$$

- 4: **end for**

- 5: Compute $\bar{u}_N = \frac{1}{N} \sum_{k=0}^{N-1} u_k$.
-

For the analysis of the mirror prox algorithm, we need the following auxiliary results. The first one is used for the mirror steps applied to the variables θ, f, h .

Lemma B.3. Let \mathcal{H} be (possibly finite-dimensional) Hilbert space and let $H \subset \mathcal{H}$ be convex and closed. Let $\tilde{h} \in H$ and $\xi, \tilde{\xi} \in \mathcal{H}^* = \mathcal{H}$, and

$$h = \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \xi \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_\eta^{h, H}(\tilde{h}, \xi), \quad (65)$$

$$\tilde{h}_+ = \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \tilde{\xi} \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_\eta^{h, H}(\tilde{h}, \tilde{\xi}). \quad (66)$$

Then, for any $\hat{h} \in H$

$$\langle h - \hat{h}, \eta \tilde{\xi} \rangle \leq \frac{1}{2} \|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 + \frac{\eta^2}{2} \|\tilde{\xi} - \xi\|_{\mathcal{H}}^2 - \frac{1}{2} \|h - \tilde{h}\|_{\mathcal{H}}^2. \quad (67)$$

Proof. By the optimality condition in (66), we have for all $\hat{h} \in H$

$$\langle \eta \tilde{\xi} - (\tilde{h} - \tilde{h}_+), \hat{h} - \tilde{h}_+ \rangle \geq 0. \quad (68)$$

Rearranging, we obtain

$$\langle \tilde{h}_+ - \hat{h}, \eta \tilde{\xi} \rangle \leq \langle \tilde{h}_+ - \hat{h}, \tilde{h} - \tilde{h}_+ \rangle = -\frac{1}{2} \|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 + \frac{1}{2} \|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2. \quad (69)$$

In the same way, by the optimality condition in (65), we have for all $\hat{h} \in H$, and, in particular for \tilde{h}_+

$$\langle \eta\xi - (\tilde{h} - h), \tilde{h}_+ - h \rangle \geq 0. \quad (70)$$

Rearranging, we obtain

$$\langle h - \tilde{h}_+, \eta\xi \rangle \leq \langle h - \tilde{h}_+, \tilde{h} - h \rangle = -\frac{1}{2}\|h - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2}\|\tilde{h}_+ - h\|_{\mathcal{H}}^2 + \frac{1}{2}\|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2. \quad (71)$$

Combining the last inequality with (69) and using the Fenchel inequality, we obtain

$$\langle h - \hat{h}, \eta\tilde{\xi} \rangle = \langle \tilde{h}_+ - \hat{h}, \eta\tilde{\xi} \rangle + \langle h - \tilde{h}_+, \eta\xi \rangle + \langle h - \tilde{h}_+, \eta(\tilde{\xi} - \xi) \rangle \quad (72)$$

$$- \frac{1}{2}\|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2}\|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 + \frac{1}{2}\|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2 \quad (73)$$

$$- \frac{1}{2}\|h - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2}\|\tilde{h}_+ - h\|_{\mathcal{H}}^2 + \frac{1}{2}\|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2 \quad (74)$$

$$+ \frac{\eta^2}{2}\|\tilde{\xi} - \xi\|_{\mathcal{H}}^2 - \frac{1}{2}\|h - \tilde{h}\|_{\mathcal{H}}^2, \quad (75)$$

which gives the result of the Lemma. \square

The second result characterizes the mirror step with respect to the measure μ .

Lemma B.4 ((Hsieh et al., 2019) [Lemma 5]). *Let $\tilde{\mu} \in \mathcal{M}$ and $\xi, \tilde{\xi} \in \mathcal{F}$, and*

$$\mu = \text{Mirr}_\eta^\mu(\tilde{\mu}, \xi), \quad (76)$$

$$\tilde{\mu}_+ = \text{Mirr}_\eta^\mu(\tilde{\mu}, \tilde{\xi}). \quad (77)$$

Then, for any $\hat{\mu} \in \mathcal{M}$

$$\langle \mu - \hat{\mu}, \eta\tilde{\xi} \rangle \leq D_\Phi(\hat{\mu}, \tilde{\mu}) - D_\Phi(\hat{\mu}, \tilde{\mu}_+) + \frac{\eta^2}{8}\|\tilde{\xi} - \xi\|_{L^\infty}^2 - 2\|\mu - \tilde{\mu}\|_{TV}^2. \quad (78)$$

The following result gives the convergence rate of Algorithm 2.

Theorem B.5. *Let Assumptions B.1, B.2 hold. Let also the stepsizes in Algorithm 2 satisfy $\eta_\theta = \eta_f = \eta_\mu = \eta_h = \frac{1}{16L}$, where L is defined in (64). Then, for any compact set $U = U_\theta \times U_f \times U_\mu \times U_h \subseteq \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$, the sequence $(\bar{\theta}_N, \bar{f}_N, \bar{\mu}_N, \bar{h}_N)$ generated by Algorithm 2 satisfies¹*

$$\begin{aligned} & \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \\ & \leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right). \end{aligned}$$

Proof. Applying Lemma B.3 to the step in θ , we obtain for any $\theta \in \Theta$ and for $k = 0, \dots, N-1$

$$\langle \theta_k - \theta, \eta_\theta F'_\theta(u_k) \rangle \leq \frac{1}{2}\|\theta - \tilde{\theta}_k\|_2^2 - \frac{1}{2}\|\theta - \tilde{\theta}_{k+1}\|_2^2 - \frac{1}{2}\|\theta_k - \tilde{\theta}_k\|_2^2 + \frac{\eta_\theta^2}{2}\|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2. \quad (79)$$

Summing these inequalities for $k = 0, \dots, N-1$, we obtain

$$\begin{aligned} & \sum_{k=0}^{N-1} \langle \theta_k - \theta, F'_\theta(u_k) \rangle \leq \frac{1}{2\eta_\theta}\|\theta - \tilde{\theta}_0\|_2^2 + \frac{R_\theta}{2\eta_\theta}, \\ & R_\theta = \sum_{k=0}^{N-1} \left(-\|\theta_k - \tilde{\theta}_k\|_2^2 + \eta_\theta^2 \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \right). \end{aligned}$$

¹There was an obvious typo in the primal-dual gap expression in the original theorem statement in the main text. We have fixed this in the following result.

In the same way, we obtain for all $f \in \mathcal{H}_f$

$$\begin{aligned} \sum_{k=0}^{N-1} \langle f_k - f, F'_f(u_k) \rangle &\leq \frac{1}{2\eta_f} \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + \frac{R_f}{2\eta_f}, \\ R_f &= \sum_{k=0}^{N-1} \left(-\|f_k - \tilde{f}_k\|_{\mathcal{H}_f}^2 + \eta_f^2 \|F'_f(u_k) - F'_f(\tilde{u}_k)\|_{\mathcal{H}_f}^2 \right) \end{aligned}$$

and for all $h \in H$

$$\begin{aligned} \sum_{k=0}^{N-1} \langle h_k - h, -F'_h(u_k) \rangle &\leq \frac{1}{2\eta_h} \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 + \frac{R_h}{2\eta_h}, \\ R_h &= \sum_{k=0}^{N-1} \left(-\|h_k - \tilde{h}_k\|_{\mathcal{H}_h}^2 + \eta_h^2 \|F'_h(u_k) - F'_h(\tilde{u}_k)\|_{\mathcal{H}_h}^2 \right). \end{aligned}$$

Finally, applying Lemma B.4 to the step in μ , we obtain for any $\mu \in \mathcal{M}$

$$\begin{aligned} \sum_{k=0}^{N-1} \langle \mu_k - \mu, -F'_\mu(u_k) \rangle &\leq \frac{1}{\eta_\mu} D_\Phi(\mu, \tilde{\mu}_0) + \frac{R_\mu}{2\eta_\mu}, \\ R_\mu &= \sum_{k=0}^{N-1} \left(-\|\mu_k - \tilde{\mu}_k\|_{TV}^2 + \eta_\mu^2 \|F'_\mu(u_k) - F'_\mu(\tilde{u}_k)\|_{L^\infty}^2 \right). \end{aligned}$$

By convexity of F in (θ, f) and concavity of F in (μ, h) , we have, for all $(\theta, f) \in \Theta \times \mathcal{H}_f$

$$\begin{aligned} \frac{1}{N} \sum_{k=0}^{N-1} F(u_k) - F(\theta, f, \bar{\mu}_N, \bar{h}_N) &\leq \frac{1}{N} \sum_{k=0}^{N-1} (F(u_k) - F(\theta, f, \mu_k, h_k)) \\ &\leq \frac{1}{N} \sum_{k=0}^{N-1} (\langle \theta_k - \theta, F'_\theta(u_k) \rangle + \langle f_k - f, F'_f(u_k) \rangle) \\ &\leq \frac{1}{2N\eta_\theta} \|\theta - \tilde{\theta}_0\|_2^2 + \frac{R_\theta}{2N\eta_\theta} + \frac{1}{2N\eta_f} \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + \frac{R_f}{2N\eta_f}. \end{aligned}$$

In the same way, we obtain that, for all $(\mu, h) \in \mathcal{M} \times H$

$$\begin{aligned} -\frac{1}{N} \sum_{k=0}^{N-1} F(u_k) + F(\bar{\theta}_N, \bar{f}_N, \mu, h) &\leq \frac{1}{N} \sum_{k=0}^{N-1} (-F(u_k) + F(\theta_k, f_k, \mu, h)) \\ &\leq \frac{1}{N} \sum_{k=0}^{N-1} (\langle \mu_k - \mu, -F'_\mu(u_k) \rangle + \langle h_k - h, -F'_h(u_k) \rangle) \\ &\leq \frac{1}{N\eta_\mu} D_\Phi(\mu, \tilde{\mu}_0) + \frac{R_\mu}{2N\eta_\mu} + \frac{1}{2N\eta_h} \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 + \frac{R_h}{2N\eta_h}. \end{aligned}$$

Combining the last two bounds, we obtain that for all $\theta \in \Theta, f \in \mathcal{H}_f, \mu \in \mathcal{M}, h \in H$ it holds that

$$\begin{aligned} &F(\bar{\theta}_N, \bar{f}_N, \mu, h) - F(\theta, f, \bar{\mu}_N, \bar{h}_N) \\ &\leq \frac{1}{2N\eta_\theta} \|\theta - \tilde{\theta}_0\|_2^2 + \frac{1}{2N\eta_f} \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + \frac{1}{N\eta_\mu} D_\Phi(\mu, \tilde{\mu}_0) + \frac{1}{2N\eta_h} \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \\ &+ \frac{1}{2N} \left(\frac{R_\theta}{\eta_\theta} + \frac{R_f}{\eta_f} + \frac{R_\mu}{\eta_\mu} + \frac{R_h}{\eta_h} \right). \end{aligned}$$

Our next goal is to show that

$$\frac{R_\theta}{\eta_\theta} + \frac{R_f}{\eta_f} + \frac{R_\mu}{\eta_\mu} + \frac{R_h}{\eta_h} \leq 0.$$

Using the Lipschitz condition in Assumption B.2, we obtain

$$\begin{aligned} R_\theta &= \sum_{k=0}^{N-1} \left(-\|\theta_k - \tilde{\theta}_k\|_2^2 + \eta_\theta^2 \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \right) \\ &\leq \sum_{k=0}^{N-1} \left(-\|\theta_k - \tilde{\theta}_k\|_2^2 + 4\eta_\theta^2 (L_{\theta\theta}^2 \|\theta_k - \tilde{\theta}_k\|_2^2 + L_{\theta f}^2 \|f_k - \tilde{f}_k\|_{\mathcal{H}_f}^2 + L_{\theta\mu}^2 \|\mu_k - \tilde{\mu}_k\|_{TV}^2 + L_{\theta h}^2 \|h_k - \tilde{h}_k\|_{\mathcal{H}_h}^2) \right). \end{aligned}$$

Combining this with the similar estimates for R_f , R_μ , R_h and rearranging the terms, we obtain

$$\begin{aligned} \frac{R_\theta}{\eta_\theta} + \frac{R_f}{\eta_f} + \frac{R_\mu}{\eta_\mu} + \frac{R_h}{\eta_h} &\leq \sum_{k=0}^{N-1} \left(\|\theta_k - \tilde{\theta}_k\|_2^2 (-1/\eta_\theta + 4(L_{\theta\theta}^2 \eta_\theta + L_{f\theta}^2 \eta_f + L_{\mu\theta}^2 \eta_\mu + L_{h\theta}^2 \eta_h)) \right. \\ &\quad + \|f_k - \tilde{f}_k\|_{\mathcal{H}_f}^2 (-1/\eta_f + 4(L_{\theta f}^2 \eta_\theta + L_{f f}^2 \eta_f + L_{\mu f}^2 \eta_\mu + L_{h f}^2 \eta_h)) \\ &\quad + \|\mu_k - \tilde{\mu}_k\|_{TV}^2 (-1/\eta_\mu + 4(L_{\theta\mu}^2 \eta_\theta + L_{f\mu}^2 \eta_f + L_{\mu\mu}^2 \eta_\mu + L_{h\mu}^2 \eta_h)) \\ &\quad \left. + \|h_k - \tilde{h}_k\|_{\mathcal{H}_h}^2 (-1/\eta_h + 4(L_{\theta h}^2 \eta_\theta + L_{f h}^2 \eta_f + L_{\mu h}^2 \eta_\mu + L_{h h}^2 \eta_h)) \right) \leq 0, \end{aligned}$$

where we used that $\eta_\theta, \eta_f, \eta_\mu, \eta_h \leq \frac{1}{16L}$ for L defined in (64).

Thus, we finally obtain that for any compact $U = U_\theta \times U_f \times U_\mu \times U_h \subset \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$

$$\begin{aligned} &\max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \\ &\leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right). \end{aligned}$$

□

B.3. Analysis in the stochastic case

To account for potential inexactness in the first-order information, we assume that instead of exact derivatives, the algorithm uses their inexact counterparts $\tilde{F}'_\theta(u)$, $\tilde{F}'_f(u)$, $\tilde{F}'_\mu(u)$, $\tilde{F}'_h(u)$, that may be random and are assumed to satisfy the following assumption.

Assumption B.6.

$$F'_\theta(u) = \mathbb{E}\tilde{F}'_\theta(u), \quad (80)$$

$$F'_f(u) = \mathbb{E}\tilde{F}'_f(u), \quad (81)$$

$$F'_\mu(u) = \mathbb{E}\tilde{F}'_\mu(u), \quad (82)$$

$$F'_h(u) = \mathbb{E}\tilde{F}'_h(u), \quad (83)$$

$$\mathbb{E}\|F'_\theta(u) - \tilde{F}'_\theta(u)\|_2^2 \leq \sigma_\theta^2, \quad (84)$$

$$\mathbb{E}\|F'_f(u) - \tilde{F}'_f(u)\|_{\mathcal{H}_f}^2 \leq \sigma_f^2, \quad (85)$$

$$\mathbb{E}\|F'_\mu(u) - \tilde{F}'_\mu(u)\|_{L^\infty}^2 \leq \sigma_\mu^2, \quad (86)$$

$$\mathbb{E}\|F'_h(u) - \tilde{F}'_h(u)\|_{\mathcal{H}_h}^2 \leq \sigma_h^2. \quad (87)$$

Theorem B.7. *Let Assumptions B.1–B.6 hold. Let also in Algorithm 2 the stochastic derivatives be used instead of the deterministic and the stepsizes satisfy $\eta_\theta = \eta_f = \eta_\mu = \eta_h = \frac{1}{16L}$, where L is defined in (64). Then, for any compact set*

$U = U_\theta \times U_f \times U_\mu \times U_h \subseteq \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$, the sequence $(\bar{\theta}_N, \bar{f}_N, \bar{\mu}_N, \bar{h}_N)$ generated by Algorithm 2 satisfies

$$\begin{aligned} & \mathbb{E} \left\{ \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \right\} \\ & \leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right) + \frac{3}{16L} (\sigma_\theta^2 + \sigma_f^2 + \sigma_\mu^2 + \sigma_h^2). \end{aligned}$$

Proof. We proceed as in the proof of Theorem B.5 changing in Algorithm 2 the exact first-order information to its inexact counterpart. In this way, we obtain the following counterpart of (79)

$$\langle \theta_k - \theta, \eta_\theta \tilde{F}'_\theta(u_k) \rangle \leq \frac{1}{2} \|\theta - \tilde{\theta}_k\|_2^2 - \frac{1}{2} \|\theta - \tilde{\theta}_{k+1}\|_2^2 - \frac{1}{2} \|\theta_k - \tilde{\theta}_k\|_2^2 + \frac{\eta_\theta^2}{2} \|\tilde{F}'_\theta(u_k) - \tilde{F}'_\theta(\tilde{u}_k)\|_2^2. \quad (88)$$

Using the inequality

$$\mathbb{E} \|\tilde{F}'_\theta(u_k) - \tilde{F}'_\theta(\tilde{u}_k)\|_2^2 \leq 3\mathbb{E} \left(\|\tilde{F}'_\theta(u_k) - F'_\theta(u_k)\|_2^2 + \|\tilde{F}'_\theta(\tilde{u}_k) - F'_\theta(\tilde{u}_k)\|_2^2 + \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \right) \quad (89)$$

$$\stackrel{\text{Assumpt. B.6}}{\leq} 6\sigma_\theta^2 + 3\mathbb{E} \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \quad (90)$$

and taking the expectation in the previous inequality, we obtain the following counterpart of (79)

$$\langle \theta_k - \theta, \eta_\theta F'_\theta(u_k) \rangle \leq \frac{1}{2} \mathbb{E} \|\theta - \tilde{\theta}_k\|_2^2 - \frac{1}{2} \mathbb{E} \|\theta - \tilde{\theta}_{k+1}\|_2^2 - \frac{1}{2} \mathbb{E} \|\theta_k - \tilde{\theta}_k\|_2^2 \quad (91)$$

$$+ \frac{3\eta_\theta^2}{2} \mathbb{E} \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 + 3\eta_\theta^2 \sigma_\theta^2. \quad (92)$$

Repeating the same steps as in the proof of Theorem B.5, we obtain that for any compact $U = U_\theta \times U_f \times U_\mu \times U_h \subset \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$

$$\begin{aligned} & \mathbb{E} \left\{ \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \right\} \\ & \leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right) + \frac{3}{16L} (\sigma_\theta^2 + \sigma_f^2 + \sigma_\mu^2 + \sigma_h^2). \end{aligned} \quad (93)$$

□

Let us denote $\sigma^2 = \sigma_\theta^2 + \sigma_f^2 + \sigma_\mu^2 + \sigma_h^2$. As we see, Theorem B.7 guarantees the same convergence rate as in the exact case, but up to some vicinity which is governed by the level of noise. In most cases, the σ^2/L term can be made of the same order $1/N$ by using mini-batching technique. Indeed, a mini-batch of size N allows to change the variance from σ^2 to σ^2/N . Yet, we note that in this case, N iterations will require the number of samples $O(N^2)$.

An alternative would be to use the information about the diameter of the set U . Indeed, assume that

$$\max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right) \leq \Omega_U^2.$$

Then, we obtain the following counterpart of the r.h.s. of (93) substituting $\eta_\theta = \eta_f = \eta_\mu = \eta_h = \eta$

$$\frac{\Omega_U^2}{2N\eta} + 3\sigma^2\eta.$$

Fixing the number of steps N and choosing

$$\eta_\theta = \eta_f = \eta_\mu = \eta_h = \eta = \min \left\{ \frac{1}{16L}, \frac{\Omega_U \sigma}{\sqrt{6N}} \right\},$$

we obtain the following result

$$\begin{aligned} & \mathbb{E} \left\{ \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \right\} \\ & \leq \max \left\{ \frac{8L\Omega_U^2}{N}, \sqrt{\frac{3\sigma^2\Omega_U^2}{2N}} \right\}. \end{aligned} \quad (94)$$

B.4. Case Study: Distributionally Robust Optimization

In this subsection we particularize the elements of Algorithm 2 for the specific DRO problem (10). We choose $\mathcal{H}_f = \mathcal{H}_h = \mathcal{H}$ to be a reproducing kernel Hilbert space with kernel k .

Our main assumptions for this problem are

- l is convex w.r.t. θ .
- $L_0 = \sup_{x, \theta} \|\nabla_{\theta} l(\theta; x)\|_2 < +\infty$.
- $\nabla_{\theta} l(\theta; x)$ is $L(x)$ -Lipschitz w.r.t. θ and $L_1 = \sup_{\mu} \mathbb{E}_{x \sim \mu} L(x)^2 < +\infty$.
- $C = \sup_x k(x, x) < +\infty$.

Clearly, then the objective F is convex in (θ, f) for fixed (μ, h) and concave in (μ, h) for fixed (θ, f) .

The Frechet derivatives of F with respect to the variables (θ, f, μ, h) are given by

$$F'_{\theta} = \mathbb{E}_{x \sim \mu} \nabla_{\theta} l(\theta; x) \quad (95)$$

$$F'_f = \int k(x, x') d\hat{\mu}(x') + \epsilon h(x) - \int k(x, x') d\mu(x') = \mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h(\cdot) - \mathbb{E}_{x \sim \mu} k(\cdot, x) \quad (96)$$

$$-F'_{\mu} = f(\cdot) - l(\theta; \cdot) \quad (97)$$

$$-F'_h = -\epsilon f(\cdot). \quad (98)$$

Since the derivative w.r.t. θ and f have the form of expectation, we can use the following stochastic counterparts. We can take a sample of X_i 's from μ to construct an unbiased stochastic derivative

$$\tilde{F}'_{\theta} = \frac{1}{N_{\theta}} \sum_{i=1}^{N_{\theta}} \nabla_{\theta} l(\theta; X_i). \quad (99)$$

Similarly, we can take a sample of X_i 's from μ and \hat{X}_i from $\hat{\mu}$ to construct an unbiased stochastic derivative

$$\tilde{F}'_f = \epsilon h(\cdot) + \frac{1}{N_f} \sum_{i=1}^{N_f} (k(\cdot, \hat{X}_i) + k(\cdot, X_i)). \quad (100)$$

The Lipschitz constants of the derivatives are estimated in the following way. The derivative F'_{θ} depends only on μ and θ . Thus, $L_{\theta f} = L_{\theta h} = 0$. Further, we have

$$\|\mathbb{E}_{x \sim \mu} \nabla_{\theta} l(\theta_1; x) - \mathbb{E}_{x \sim \mu} \nabla_{\theta} l(\theta_2; x)\|_2 \leq \mathbb{E}_{x \sim \mu} L(x) \|\theta_1 - \theta_2\|_2 \quad (101)$$

and $L_{\theta\theta} = L_1$.

$$\|\mathbb{E}_{x \sim \mu_1} \nabla_{\theta} l(\theta; x) - \mathbb{E}_{x \sim \mu_2} \nabla_{\theta} l(\theta; x)\|_2 \leq L_0 \|\mu_1 - \mu_2\|_{TV} \quad (102)$$

and $L_{\theta\mu} = L_0$.

The derivative F'_f depends only on μ and h . Thus, $L_{ff} = L_{f\theta} = 0$. Further, we have

$$\|\mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h_1(\cdot) + \mathbb{E}_{x \sim \mu_1} k(\cdot, x) - (\mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h_2(\cdot) + \mathbb{E}_{x \sim \mu_2} k(\cdot, x))\|_{\mathcal{H}} \quad (103)$$

$$\leq \epsilon \|h_1 - h_2\|_{\mathcal{H}} + \|\mathbb{E}_{x \sim \mu_1} k(\cdot, x) - \mathbb{E}_{x \sim \mu_2} k(\cdot, x)\|_{\mathcal{H}} \leq \epsilon \|h_1 - h_2\|_{\mathcal{H}} + \sqrt{C} \|\mu_1 - \mu_2\|_{TV}, \quad (104)$$

i.e., $L_{f\mu} = \sqrt{C}$, $L_{fh} = \epsilon$. Here we used that

$$\|\mathbb{E}_{x \sim \mu_1} k(\cdot, x) - \mathbb{E}_{x \sim \mu_2} k(\cdot, x)\|_{\mathcal{H}} \leq \sqrt{C} \|\mu_1 - \mu_2\|_{TV}.$$

The derivative F'_μ depends only on f and θ . Thus, $L_{\mu\mu} = L_{\mu h} = 0$. Further, we have

$$\| -f_1(\cdot) + l(\theta_1; \cdot) - (-f_2(\cdot) + l(\theta_2; \cdot)) \|_{L^\infty} \quad (105)$$

$$\leq \sqrt{C} \|f_1 - f_2\|_{\mathcal{H}} + L_0 \|\theta_1 - \theta_2\|_{\mathcal{H}}, \quad (106)$$

i.e., $L_{f\mu} = \sqrt{C}$, $L_{f\theta} = L_0$. Here we used that

$$\begin{aligned} \|f_2(\cdot) - f_1(\cdot)\|_{L^\infty} &= \sup_x |f_2(x) - f_1(x)| = \sup_x \langle f_2 - f_1, \phi(x) \rangle_{\mathcal{H}} \\ &\leq \|f_2 - f_1\|_{\mathcal{H}} \cdot \sup_x \|\phi(x)\|_{\mathcal{H}} \leq \sqrt{C} \cdot \|f_2 - f_1\|_{\mathcal{H}}. \end{aligned} \quad (107)$$

Finally, the derivative F'_h depends only on f . Thus, $L_{h\theta} = L_{h\mu} = L_{hh} = 0$. Further, we have

$$\| -\epsilon f_1(\cdot) - (-\epsilon f_2(\cdot)) \|_{\mathcal{H}} \leq \epsilon \|f_1 - f_2\|_{\mathcal{H}}. \quad (108)$$

Thus, $L_{hf} = \epsilon$.

As we see, our main assumptions in Theorem B.5 hold for the DRO problem (10). Moreover, stochastic derivatives in (99) and (100) satisfy Assumption B.6. Indeed, we have

$$\mathbb{E}_{X \sim \mu} \|\mathbb{E}_{x \sim \mu} \nabla_\theta l(\theta; x) - \nabla_\theta l(\theta; X)\|_2^2 \leq \mathbb{E}_{X \sim \mu} \|\nabla_\theta l(\theta; X)\|_2^2 \leq L_0^2, \quad (109)$$

$$\mathbb{E}_{\hat{X} \sim \hat{\mu}, X \sim \mu} \left\| \epsilon h(\cdot) + k(\cdot, \hat{X}) + k(\cdot, X) - (\mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h(\cdot) - \mathbb{E}_{x \sim \mu} k(\cdot, x)) \right\|_{\mathcal{H}}^2 \quad (110)$$

$$\leq 2\mathbb{E}_{X \sim \mu} \|k(\cdot, X)\|_{\mathcal{H}}^2 + 2\mathbb{E}_{\hat{X} \sim \hat{\mu}} \|k(\cdot, \hat{X})\|_{\mathcal{H}}^2 \leq 4C. \quad (111)$$

This allows us to apply also Theorem B.7 to the DRO problem (10). This proves Corollary 5.1.